Algorithmic bias? A study of data-based discrimination in the serving of ads in social media

Anja Lambrecht and Catherine Tucker

Research Question

What may make an ad serving algorithm appear biased?

Motivation

- Privacy debate has moved to a question of privacy harms:
- Papers in CS have documented empirical pattern of apparently discriminatory ad serving behavior (Sweeney, 2013; Datta et al., 2015)
- But they are not focused on understanding why

What we do

- Field Test data on STEM ad across 190 countries
  - Set up as gender neutral
  - But shown to men more than women

Why apparent algorithmic-bias happens

- Not because of
  - Click propensity
  - Media usage
  - Underlying sexism
- Evidence that young women are valuable demographic and other advertiser bids crowd out intentionally gender neutral advertisers

Why does this matter?

- First paper to explore the why of apparent algorithmic-bias
- We find that apparent algorithmic bias may not be intentional but instead the result of completely separate advertiser actions
- Emphasizes that privacy online is not an individual issue. Instead it may be a complex mass of intertwined decisions.

The FTC is worried about algorithmic transparency, and you should be too

Policy Implications

- Not much support in our findings for ‘Algorithmic Transparency’ being a solution
  - Perhaps auditing algorithmic outcomes is a better approach.
- If regulating privacy in online advertising is hard, regulating the potential for algorithmic discrimination or bias may be even harder

Outline

Origin of the Test

Desktop Right Column

STEM Careers
Information about STEM Careers

Figure: Sample Ad

This was a very straightforward field test

- All that varied was the country it was targeted at
- 191 countries
- Ensured that in each country the ad was shown at least to 5000 people

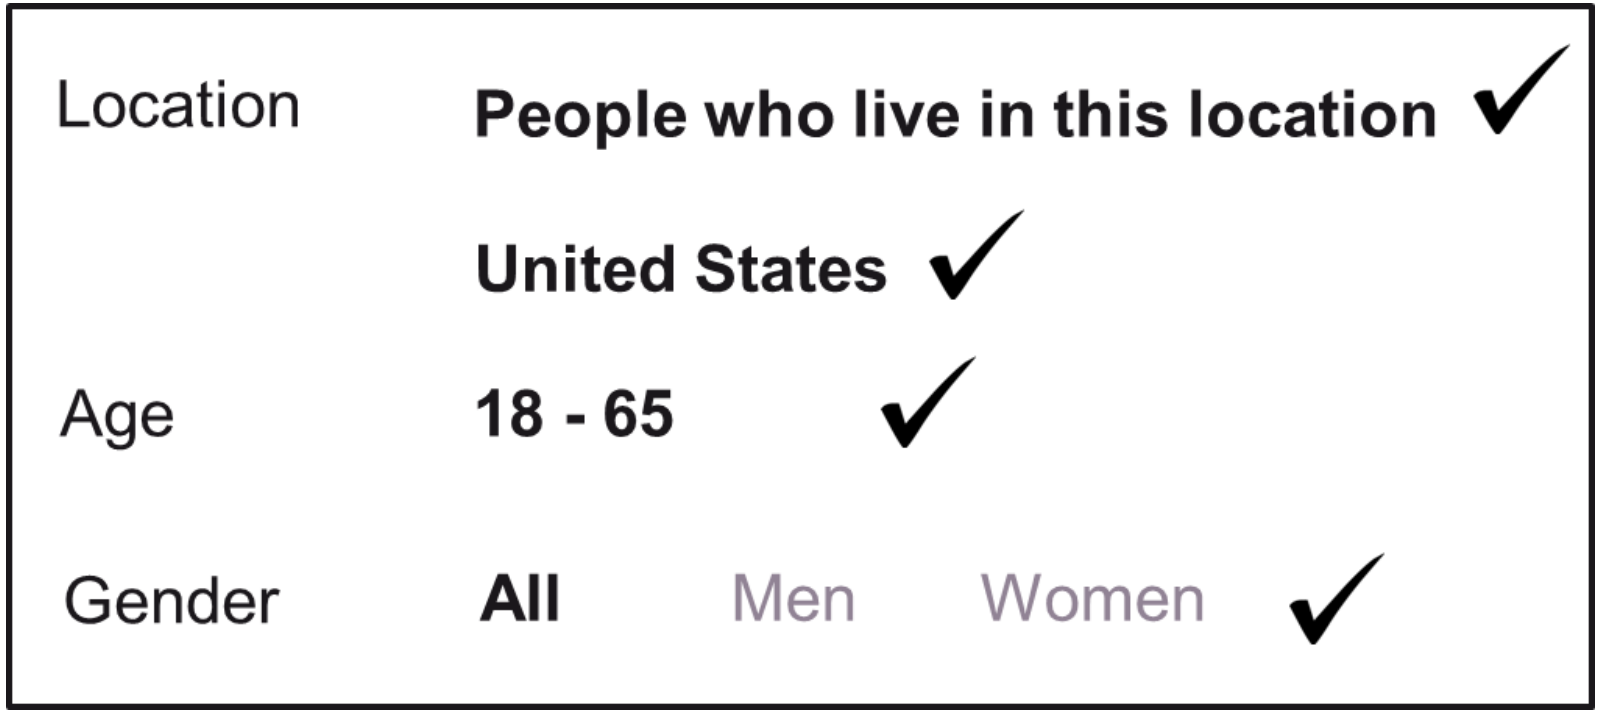

Figure: Ad Targeting Settings - Ad intended to be shown to both men and women aged 18-65.

	Mean	Std Dev	Min	Max
Impressions	1911.8	2321.4	0	24980
Clicks	3.00	4.52	0	42
Unique Clicks	2.78	4.15	0	40
CPC	0.085	0.090	0	0.66
CPM	0.18	0.30	0	3.85
Reach	615.6	850.7	0	13436
Frequency	4.38	4.32	1	53
Click Rate	0.15	0.17	0	1.52
Reach Rate	0.0064	0.013	0	0.25
Female	0.50	0.50	0	1
(mean) femalelaborpart	74.4	16.3	18.7	103.6
(mean) femaleprimary	103.4	17.0	20.8	174.8
(mean) femaleequality	3.31	0.58	1.50	4.50

Table: Summary statistics

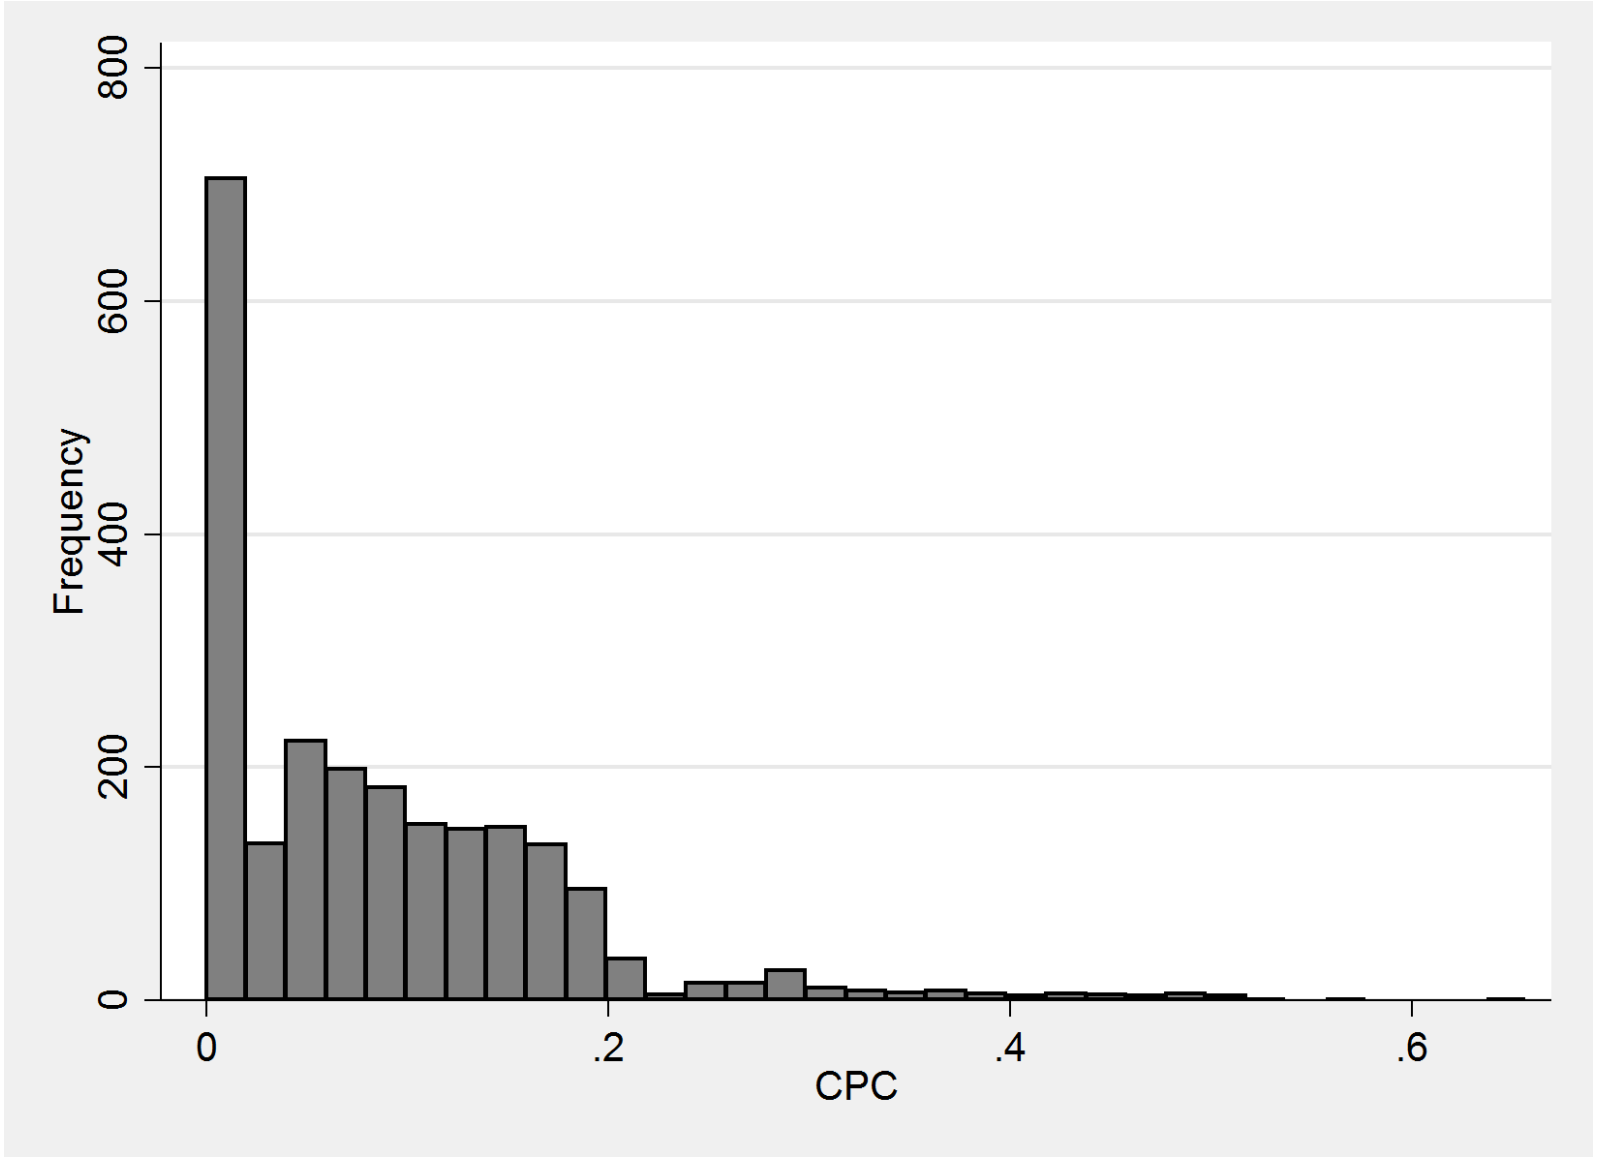

Figure: Histogram of average cost per country

Outline

Really, this paper doesn’t need any complex analysis

Table: Raw Data reported

Age Group	Male Impr.	Female Impr.	Male Clicks	Female Clicks
Age18-24	746719	649590	1156	1171
Age25-34	662996	495996	873	758
Age35-44	412457	283596	501	480
Age45-54	307701	224809	413	414
Age55-64	209608	176454	320	363
Age 65+	192317	153470	307	321

Table: Raw Data Reported as an Average per Country

Age Group	Male Impr.	Female Impr.	Male Clicks	Female Clicks
Age18-24	3909	3401	6	6
Age25-34	3471	2597	5	4
Age35-44	2159	1485	3	3
Age45-54	1611	1177	2	2
Age55-64	1097	924	2	2
Age 65+	1007	808	2	2

Three obvious patterns in the data

- Men see more impressions of the ad than women.
- Particularly in younger ad cohorts
- Clicks appear similar

Outline

Really, this paper doesn’t need any complex analysis

For campaign i and demographic group j in country k on day t, the number of times an ad is displayed is modeled as a function of:

$$
\begin{aligned}
AdDisplay_{ijkt} = & \\
& + \quad \beta_1 Female_j \\
& + \quad \beta_2 Age_j \\
& + \quad \beta_3 Female_j \times Age_j \\
& + \quad \alpha_k + \epsilon_{jk} \qquad (1)
\end{aligned}
$$

Table: Women Are Shown Fewer Ads Than Men

	(1) Impressions	(2) Impressions	(3) Reach	(4) Reach	(5) Frequency	(6) Frequency
Female	-479.3***	-209.7***	-228.1***	-98.97***	0.729***	1.276***
	(97.09)	(44.26)	(35.45)	(20.44)	(0.150)	(0.305)
Female × Age18-24		-298.8		-234.3**		-0.523
		(193.1)		(75.83)		(0.268)
Female × Age25-34		-664.6***		-302.2***		-0.630*
		(154.4)		(48.64)		(0.272)
Female × Age35-44		-464.9***		-159.9***		-0.900***
		(110.5)		(31.26)		(0.246)
Female × Age45-54		-224.2**		-97.25***		-0.903**
		(69.94)		(24.70)		(0.300)
Female × Age55-64		36.16		18.93		-0.326
		(39.58)		(14.33)		(0.412)
Age18-24	2753.6***	2902.6***	909.5***	1026.5***	-0.473*	-0.212
	(248.0)	(284.3)	(108.5)	(131.2)	(0.207)	(0.174)
Age25-34	2132.4***	2464.3***	561.4***	712.3***	-0.683***	-0.369*
	(204.4)	(236.5)	(67.32)	(83.38)	(0.163)	(0.143)
Age35-44	920.5***	1152.6***	197.4***	277.2***	-0.556***	-0.107
	(117.4)	(135.2)	(40.61)	(47.39)	(0.144)	(0.167)
Age45-54	492.4***	604.1***	99.08**	147.5***	-0.471***	-0.0198
	(84.60)	(85.93)	(31.03)	(35.27)	(0.108)	(0.167)
Age55-64	109.0*	90.53	16.56	6.911	0.0107	0.173
	(51.37)	(52.72)	(18.93)	(19.70)	(0.182)	(0.147)
Country Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2291	2291	2291	2291	2291	2291
R-Squared	0.485	0.488	0.442	0.446	0.776	0.778

Ordinary Least Squares Estimates. Dependent variable as shown. Omitted demographic groups are those aged 65+ and men. Robust standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Do our results reflect the fact that women were less likely to click on the ad?

Table: If They See The Ad, Women Are More Likely To Click Than Men

	(1) Clicks	(2) Unique Clicks	(3) Click Rate	(4) Reach Rate	(5) Clicks	(6) Unique Clicks	(7) Click Rate	(8) Reach Rate
Female	0.221***	0.303***	0.0362***	0.00280***	0.264**	0.399***	0.0425	0.00366*
	(0.0271)	(0.0290)	(0.00713)	(0.000599)	(0.0932)	(0.0875)	(0.0233)	(0.00177)
Female × Age18-24					-0.137	-0.166	-0.0156	-0.00107
					(0.0975)	(0.0956)	(0.0265)	(0.00164)
Female × Age25-34					-0.0899	-0.135	-0.0254	-0.00223
					(0.113)	(0.109)	(0.0283)	(0.00209)
Female × Age35-44					0.0822	-0.0289	-0.0136	-0.00244
					(0.113)	(0.109)	(0.0273)	(0.00196)
Female × Age45-54					0.0633	0.000689	-0.00486	-0.00180
					(0.119)	(0.117)	(0.0288)	(0.00178)
Female × Age55-64					0.0465	-0.0573	0.0221	0.00238
					(0.136)	(0.129)	(0.0308)	(0.00221)
Age18-24	-0.175**	-0.214***	-0.0216	-0.00117	-0.105	-0.129	-0.0138	-0.000637
	(0.0576)	(0.0557)	(0.0139)	(0.000825)	(0.0731)	(0.0704)	(0.0152)	(0.000585)
Age25-34	-0.375***	-0.460***	-0.0500***	-0.00271**	-0.332***	-0.394***	-0.0373*	-0.00160*
	(0.0593)	(0.0572)	(0.0127)	(0.000850)	(0.0823)	(0.0785)	(0.0180)	(0.000680)
Age35-44	-0.341***	-0.409***	-0.0493***	-0.00189*	-0.379***	-0.392***	-0.0425*	-0.000668
	(0.0712)	(0.0657)	(0.0133)	(0.000904)	(0.0902)	(0.0839)	(0.0174)	(0.00112)
Age45-54	-0.190**	-0.222***	-0.0288*	-0.00166	-0.220*	-0.220**	-0.0264	-0.000764
	(0.0613)	(0.0605)	(0.0123)	(0.000865)	(0.0865)	(0.0843)	(0.0158)	(0.000680)
Age55-64	-0.0186	-0.0199	-0.00190	0.00149	-0.0426	0.00913	-0.0129	0.000296
	(0.0682)	(0.0666)	(0.0149)	(0.000912)	(0.0955)	(0.0879)	(0.0167)	(0.000863)
Country Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4515014	1453890	2291	2291	4515014	1453890	2291	2291
Log-Likelihood	-52298.6	-40388.3	1055.8	7193.9	-52291.8	-40384.6	1058.5	7201.1
R-Squared			0.173	0.314			0.175	0.318

Aggregate Logit Estimates in Columns (1)-(2) and (5)-(6). Ordinary Least Squares Estimates in Columns (3)-(4) and (7)-(8). In Columns (2), (4), (6) and (8) the population variable is ad reach. In Columns (1), (3), (5), and (7) the population variable is ad impressions. The dependent variable is whether someone who was exposed to an ad clicked. Omitted demographic groups are those aged 65+ and men. Robust standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Do women spend less time on social media?

- No.
- At least every piece of recorded data says no.

Do our results reflect cultural prejudice or labor market conditions for women?

Table: Women Being Exposed To Fewer Ads Than Men Is Not Driven Entirely By Underlying Gender Disparity In Labor Market Conditions In That Country

	(1) Reach	(2) Reach	(3) Reach
Female	-326.7***	-257.3***	-324.8***
	(91.61)	(45.34)	(56.52)
Female × High % Female labor part=1	58.72		
	(100.9)		
Female × High % Female primary=1		-64.69	
		(101.0)	
Female × High Female Equality Index (CPIA)=1			140.6
			(162.3)
Age18-24	1035.3***	1007.0***	1057.3***
	(149.6)	(149.0)	(150.5)
Age25-34	620.7***	610.6***	1181.9***
	(96.55)	(95.92)	(106.1)
Age35-44	177.4**	173.1**	460.9***
	(58.79)	(58.20)	(42.14)
Age45-54	64.55	56.19	150.9***
	(45.13)	(44.42)	(32.05)
Age55-64	-12.99	-17.90	-42.40
	(27.34)	(26.89)	(27.98)
Country Controls	Yes	Yes	Yes
Observations	1500	1512	588
Log-Likelihood	-11998.5	-12091.7	-4485.8
R-Squared	0.417	0.422	0.601

Ordinary Least Squares Estimates. Dependent variable is whether someone is exposed to an ad. Omitted demographic groups are those aged 65+ and men. Robust standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Do our results simply reflect competitive spillovers?

Does price matter?

Across all campaigns, the average cost per click was nearly identical for men and women ($0.09)

But maybe we just were not bidding high enough to reach women. So we went out and collected some more data.

	Mean	Std Dev	Min	Max
Avg Suggested Bid	0.45	0.66	0.010	15.7
Min Suggested Bid	0.19	0.31	0.010	4
Max Suggested Bid	0.77	1.32	0.017	43
Female	0.50	0.50	0	1

Table: Summary statistics

Table: In General, Women Are More Expensive To Advertise To On Social Media And The Competitive Spillover From Other Advertisers' Decisions May Explain Our Finding

	(1) Avg Suggested Bid	(2) Avg Suggested Bid	(3) Min Suggested Bid	(4) Max Suggested Bid
Female	-0.0464	0.0525*	-0.0139	-0.0157
	(0.0378)	(0.0247)	(0.0294)	(0.0404)
Female × Age18-24	0.0648^{+}		0.0242	-0.221
	(0.0376)		(0.0296)	(0.282)
Female × Age25-34	0.174^{+}		0.0393	0.103*
	(0.0935)		(0.0295)	(0.0436)
Female × Age35-44	0.150***		0.0683*	0.191***
	(0.0429)		(0.0296)	(0.0481)
Female × Age45-54	0.0751		0.0235	0.128^{+}
	(0.0544)		(0.0387)	(0.0751)
Female × Age55+	0.129**		0.0496	0.193***
	(0.0445)		(0.0346)	(0.0546)
Age18-24	-0.0421	-0.0100	-0.0421	0.342
	(0.0405)	(0.0282)	(0.0310)	(0.283)
Age25-34	-0.0105	0.0763	-0.0415	0.118*
	(0.0406)	(0.0519)	(0.0310)	(0.0495)
Age35-44	-0.000557	0.0740*	-0.0477	0.173**
	(0.0444)	(0.0364)	(0.0325)	(0.0610)
Age45-54	0.0216	0.0589	-0.0268	0.229**
	(0.0557)	(0.0405)	(0.0362)	(0.0817)
Age55+	-0.0446	0.0198	-0.0551	0.102^{+}
	(0.0435)	(0.0347)	(0.0335)	(0.0591)
Country Controls	Yes	Yes	Yes	Yes
Observations	2096	2096	1916	1915
Log-Likelihood	-1215.0	-1219.8	637.1	-2745.5
R-Squared	0.571	0.569	0.679	0.409

Ordinary Least Squares Estimates. Dependent variable is average suggested bid in the Columns (1)-(3), minimum suggested bid in Column (4) and maximum suggested bid in Column (5). Omitted demographic groups are those aged between 13-17 and those of the male gender. Robust standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Why Are Women Such a Prized Demographic?

To investigate this, we looked at additional data about the purchasing of consumer items as a result of a social media campaign.

Table: Younger women may be a valuable demographic as they appear more likely to convert conditional on clicking an ad

	Clicks out of impressions	Add-to-cart out of clicks	Add-to-cart out of impressions
	(1)	(2)	(3)
	Clicks	Add to Cart	Add to Cart
Female	-0.0522***	-0.0231	-0.0979
	(0.0152)	(0.186)	(0.185)
Age Group 18-24	-0.795***	-0.528	-1.392**
	(0.0379)	(0.558)	(0.548)
Age Group 25-35	-0.533***	-0.149	-0.742***
	(0.0194)	(0.265)	(0.264)
Age Group 35-44	-0.244***	-0.168	-0.430**
	(0.0155)	(0.202)	(0.201)
Female $\times$ Age Group 18-24	0.408***	1.078*	1.553***
	(0.0399)	(0.575)	(0.566)
Female $\times$ Age Group 25-35	-0.0602**	0.701**	0.709**
	(0.0272)	(0.326)	(0.324)
Female $\times$ Age Group 35-44	-0.000403	0.509*	0.508*
	(0.0220)	(0.264)	(0.263)
Week Controls	Yes	Yes	Yes
Day of week controls	Yes	Yes	Yes
Product Controls	Yes	Yes	Yes
Observations	127617816	67501	127605845
Log-Likelihood	-574304.1	-3339.4	-7802.1

Aggregate logit estimates. Dependent variable as listed. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Omitted demographic groups are men and those aged 45+.

Outline

Limitations

- Single field test.
- Descriptive paper
- Just look at gender
- Big (non-economist) questions are not tackled - Should we think of this as bias? Should we think of this as discrimination?

Punchline

- Cross-national field test suggests that an ad which is intended to be gender-neutral may not be allocated in a gender-neutral way by an ad-serving algorithm
- We show that women are shown fewer STEM ads than men NOT because of an algorithm responding to click behavior or local prejudice
- But instead because women's desirability as a demographic and consequent high price means that an algorithm trained to be cost effective avoids showing ads to them.
- Apparent algorithmic bias may be an unintentional consequence of external behavior

Implications for Practice

- Managers can't assume an algorithm will neutrally deliver ads.
- In our case, can be easily solved by managing two separate campaigns for men and women and paying more for women.
- But what about cases where the algorithm does not neutrally distribute ads with respect to harder-to-address factors such as economic marginalization or race?

Implications for Policy

- Difficult to see how algorithmic transparency would help here?
- Emphasizes the need for nuance in algorithmic auditing policy

Thank you!

alambrecht@london.edu cetucker@mit.edu

Datta, A., M. C. Tschantz, and A. Datta (2015). Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies 2015*(1), 92–112.

Sweeney, L. (2013). Discrimination in online ad delivery. *ACMQueue 11*(3), 10.